

Learning and Forecasting Opinion Dynamics in Social Networks

Abir De¹, Isabel Valera², Niloy Ganguly¹, Sourangshu Bhattacharya¹, and Manuel Gomez-Rodriguez²

¹IIT Kharagpur, {abir.de, niloy, sourangshu}@cse.iitkgp.ernet.in

²Max Plank Institute for Software Systems, {ivalera, manuelgr}@mpi-sws.org

Abstract

Social media and social networking sites have become a global pinboard for exposition and discussion of news, topics, and ideas, where social media users often update their opinions about a particular topic by learning from the opinions shared by their friends. In this context, can we learn a data-driven model of opinion dynamics that is able to accurately forecast users' opinions? In this paper, we introduce SLANT, a probabilistic modeling framework of opinion dynamics, which represents users' opinions over time by means of marked jump diffusion stochastic differential equations, and allows for efficient model simulation and parameter estimation from historical fine grained event data. We then leverage our framework to derive a set of efficient predictive formulas for opinion forecasting and identify conditions under which opinions converge to a steady state. Experiments on data gathered from Twitter show that our model provides a good fit to the data and our formulas achieve more accurate forecasting than alternatives.

1 Introduction

Social media and social networking sites are increasingly used by people to express their opinions, give their “hot takes”, on the latest breaking news, political issues, sports events, and new products. As a consequence, there has been an increasing interest on leveraging social media and social networking sites to sense and forecast *opinions*, as well as understand *opinion dynamics*. For example, political parties routinely use social media to sense people’s opinion about their political discourse¹; quantitative investment firms measure investor sentiment and trade using social media [17]; and, corporations leverage brand sentiment, estimated from users’ posts, likes and shares in social media and social networking sites, to design their marketing campaigns². In this context, multiple methods for sensing opinions, typically based on sentiment analysis [20], have been proposed in recent years. However, methods for accurately forecasting opinions are still scarce [6, 7, 18], despite the extensive literature on theoretical models of opinion dynamics [5, 8].

In this paper, we develop a novel modeling framework of opinion dynamics in social media and social networking sites, SLANT³, which allows for accurate forecasting of individual users’ opinions. The proposed framework is based on two simple intuitive ideas: i) users’ opinions are *hidden* until they decide to *share* it with their friends (or neighbors); and, ii) users may update their opinions about a particular topic by learning from the opinions *shared* by their friends. While the latter is one of the main underlying premises used by many well-known theoretical models of opinion dynamics [5, 8, 21], the former has been ignored by models of opinion dynamics, despite its relevance on closely related processes such as information diffusion [11].

More in detail, our proposed model represents users’ *latent* opinions as continuous-time stochastic processes driven by a set of marked jump stochastic differential equations (SDEs) [13]. Such construction allows each user’s latent opinion to be modulated over time by the opinions asynchronously *expressed* by her neighbors as *sentiment* messages. Here, every time a user expresses an opinion by posting a sentiment message, she *reveals* a noisy estimate of her current latent opinion. Then, we exploit a key property of our model, the Markov property, to develop:

¹<http://www.nytimes.com/2012/10/08/technology/campaigns-use-social-media-to-lure-younger-voters.html>

²<http://www.nytimes.com/2012/07/31/technology/facebook-twitter-and-foursquare-as-corporate-focus-groups.html>

³Slant is a particular point of view from which something is seen or presented.

- I. An efficient estimation procedure to find the parameters that maximize the likelihood of a set of (millions of) sentiment messages via convex programming.
- II. A scalable simulation procedure to sample millions of sentiment messages from the proposed model in a matter of minutes.
- III. A set of novel predictive formulas for efficient and accurate opinion forecasting, which can also be used to identify conditions under which opinions converge to a steady state of consensus or polarization.

Finally, we experiment on both synthetic and real data gathered from Twitter and show that our model provides a good fit to the data and our predictive formulas achieve more accurate opinion forecasting than several alternatives [6, 7, 8, 14, 25].

Related work. There is an extensive line of work on theoretical models of opinion dynamics and opinion formation [3, 5, 8, 14, 16, 25]. However, previous models typically share the following limitations: (i) they do not distinguish between latent opinion and sentiment (or expressed opinion), which is a noisy observation of the opinion (*e.g.*, thumbs up/down, text sentiment); (ii) they consider users' opinions to be updated synchronously in discrete time, however, opinions may be updated asynchronously following complex temporal patterns [11]; (iii) the model parameters are difficult to learn from real fine-grained data and instead are set arbitrarily, as a consequence, they provide inaccurate fine-grained predictions; and, (iv) they focus on analyzing only the steady state of the users' opinions, neglecting the transient behavior of real opinion dynamics which allows for opinion forecasting methods. More recently, there have been some efforts on designing models that overcome some of the above limitations and provide more accurate predictions [6, 7]. However, they do not distinguish between opinion and sentiment and still consider opinions to be updated synchronously in discrete time. Our modeling framework addresses the above limitations and, by doing so, achieves more accurate opinion forecasting than alternatives.

2 Proposed model

In this section, we first formulate our model of opinion dynamics, starting from the data it is designed for, and then introduce efficient methods for model parameter estimation and model simulation.

Opinions data. Given a directed social network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, we record each message as $e := (u, m, t)$, where the triplet means that the user $u \in \mathcal{V}$ posted a message with sentiment m at time t . Given a collection of messages $\{e_1 = (u_1, m_1, t_1), \dots, e_n = (u_n, m_n, t_n)\}$, the history $\mathcal{H}_u(t)$ gathers all messages posted by user u up to but not including time t , *i.e.*,

$$\mathcal{H}_u(t) = \{e_i = (u_i, m_i, t_i) \mid u_i = u \text{ and } t_i < t\}, \quad (1)$$

and $\mathcal{H}(t) := \cup_{u \in \mathcal{V}} \mathcal{H}_u(t)$ denotes the entire history of messages up to but not including time t .

Generative process. We represent users' latent opinions as a multidimensional stochastic process $\mathbf{x}^*(t)$, in which the u -th entry, $x_u^*(t) \in \mathbb{R}$, represents the opinion of user u at time t and the sign $*$ means that it may depend on the history $\mathcal{H}(t)$. Then, every time a user u posts a message at time t , we draw its sentiment m from a sentiment distribution $p(m|x_u^*(t))$. Here, we can also think of the sentiment m of each message as samples from a noisy stochastic process $m_u(t) \sim p(m_u(t)|x_u^*(t))$. Further, we represent the message times by a set of counting processes. In particular, we denote the set of counting processes as a vector $\mathbf{N}(t)$, in which the u -th entry, $N_u(t) \in \{0\} \cup \mathbb{Z}^+$, counts the number of sentiment messages user u posted up to but not including time t . Then, we can characterize the message rate of the users using their corresponding conditional intensities as

$$\mathbb{E}[d\mathbf{N}(t) \mid \mathcal{H}(t)] = \boldsymbol{\lambda}^*(t) dt, \quad (2)$$

where $d\mathbf{N}(t) := (dN_u(t))_{u \in \mathcal{V}}$ denotes the number of messages per user in the window $[t, t + dt)$ and $\boldsymbol{\lambda}^*(t) := (\lambda_u^*(t))_{u \in \mathcal{V}}$ denotes the associated user intensities, which may depend on the history $\mathcal{H}(t)$. We denote the set of user that u follows by $\mathcal{N}(u)$. Next, we specify the the intensity functions $\boldsymbol{\lambda}^*(t)$, the dynamics of the users' opinions $\mathbf{x}^*(t)$, and the sentiment distribution $p(m|x_u^*(t))$.

Intensity for messages. There is a wide variety of message intensity functions one can choose from to model the users' intensity $\boldsymbol{\lambda}^*(t)$ [1]. In this work, we consider two of the most popular functional forms used in the growing literature on social activity modeling using point processes [9, 23]:

- I. **Poisson process.** The intensity is assumed to be independent of the history $\mathcal{H}(t)$ and constant, *i.e.*, $\lambda_u^*(t) = \mu_u$.

II. Multivariate Hawkes processes. The intensity captures a mutual excitation phenomena between message events and depends on the whole history of message events $\cup_{v \in \{u \cup \mathcal{N}(u)\}} \mathcal{H}_v(t)$ before t :

$$\lambda_u^*(t) = \mu_u + \sum_{v \in u \cup \mathcal{N}(u)} b_{vu} \sum_{e_i \in \mathcal{H}_v(t)} \kappa(t - t_i) = \mu_u + \sum_{v \in u \cup \mathcal{N}(u)} b_{vu} (\kappa(t) \star dN_v(t)), \quad (3)$$

where the first term, $\mu_u \geq 0$, models the publication of messages by user u on her own initiative, and the second term, with $b_{vu} \geq 0$, models the publication of additional messages by user u due to the influence that previous messages posted by the users she follows have on her intensity. Here, $\kappa(t) = e^{\nu t}$ is an exponential triggering kernel modeling the decay of influence of the past events over time and \star denotes the convolution operation.

In both cases, the couple $(\mathbf{N}(t), \lambda^*(t))$ is a Markov process, *i.e.*, future states of the process (conditional on past and present states) depends only upon the present state, and we can express the users' intensity more compactly using the following jump stochastic differential equation (SDE):

$$d\lambda^*(t) = \nu(\mu - \lambda^*(t))dt + \mathbf{B}d\mathbf{N}(t),$$

where the initial condition is $\lambda^*(0) = \mu$. The Markov property will become important later.

Stochastic process for opinion. The opinion $x_u^*(t)$ of a user u at time t adopts the following form:

$$x_u^*(t) = \alpha_u + \sum_{v \in \mathcal{N}(u)} a_{vu} \sum_{e_i \in \mathcal{H}_v(t)} m_i g(t - t_i) = \alpha_u + \sum_{v \in \mathcal{N}(u)} a_{vu} (g(t) \star (m_v(t) dN_v(t))), \quad (4)$$

where the first term, $\alpha_u \in \mathbb{R}$, models the original opinion a user u starts with, the second term, with $a_{vu} \in \mathbb{R}$, models updates in user u 's opinion due to the influence that previous messages with opinions m_i posted by the users u follows has on her opinion. Here, $g(t) = e^{-\omega t}$ denotes an exponential triggering kernel, which models the decay of influence over time.

Under this form, the resulting opinion dynamics are Markovian and can be compactly represented by a set of coupled marked jumped stochastic differential equations (proven in Appendix A):

Proposition 1. *The tuple $(\mathbf{x}^*(t), \lambda^*(t), \mathbf{N}(t))$ is a Markov process, whose dynamics are defined by the following marked jumped stochastic differential equations (SDE):*

$$d\mathbf{x}^*(t) = \omega(\alpha - \mathbf{x}^*(t))dt + \mathbf{A}(\mathbf{m}(t) \odot d\mathbf{N}(t)) \quad (5)$$

$$d\lambda^*(t) = \nu(\mu - \lambda^*(t))dt + \mathbf{B}d\mathbf{N}(t) \quad (6)$$

where the initial conditions are $\lambda^*(0) = \mu$ and $\mathbf{x}^*(0) = \alpha$, the marks are the sentiment messages $\mathbf{m}(t) = (m_u(t))_{u \in \mathcal{V}}$, with $m_u(t) \sim p(m|x_u^*(t))$, and the sign \odot denotes pointwise product.

The above mentioned Markov property will be the key to the design of efficient model parameter estimation and model simulation algorithms.

Sentiment distribution. The particular choice of sentiment distribution $p(m|x_u^*(t))$ depends on the recorded marks. For example, one may consider:

I. **Gaussian Distribution** The sentiment is assumed to be a real random variable $m \in \mathbb{R}$, *i.e.*, $p(m|x_u(t)) = \mathcal{N}(x_u(t), \sigma_u)$. This fits well scenarios in which sentiment is extracted from text using sentiment analysis [12].

II. **Logistic.** The sentiment is assumed to be a binary random variable $m \in \{-1, 1\}$, *i.e.*, $p(m|x_u(t)) = 1/(1 + \exp(-m \cdot x_u(t)))$. This fits well scenarios in which sentiment is measured by means of up votes, down votes or likes.

Our model estimation method can be easily adapted to any log-concave sentiment distribution. However, in the remainder of the paper, we consider the Gaussian distribution since, in our experiments, sentiment is extracted from text using sentiment analysis.

2.1 Model parameter estimation

Given a collection of messages $\mathcal{H}(T) = \{(u_i, m_i, t_i)\}$ recorded during a time period $[0, T]$ in a social network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, we can find the optimal parameters α , μ , \mathbf{A} and \mathbf{B} by solving a maximum likelihood estimation (MLE) problem⁴. To do so, it is easy to show that the log-likelihood of the messages is given by

⁴Here, if one decides to model the message intensities with a Poisson process, $\mathbf{B} = 0$.

$$\mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\mu}, \mathbf{A}, \mathbf{B}) = \underbrace{\sum_{e_i \in \mathcal{H}(T)} \log p(m_i | x_{u_i}^*(t_i))}_{\text{message sentiments}} + \underbrace{\sum_{e_i \in \mathcal{H}(T)} \log \lambda_{u_i}^*(t_i) - \sum_{u \in \mathcal{V}} \int_0^T \lambda_u^*(\tau) d\tau}_{\text{message times}}. \quad (7)$$

Then, we can find the optimal parameters $(\boldsymbol{\alpha}, \boldsymbol{\mu}, \mathbf{A}, \mathbf{B})$ using MLE as

$$\underset{\boldsymbol{\alpha}, \boldsymbol{\mu} \geq 0, \mathbf{A}, \mathbf{B} \geq 0}{\text{maximize}} \quad \mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\mu}, \mathbf{A}, \mathbf{B}). \quad (8)$$

Note that, as long as the sentiment distributions are log-concave, the MLE problem above is concave and thus can be solved efficiently. Moreover, the problem decomposes in $2|\mathcal{V}|$ independent subproblems, two per user u , since the first term in Eq. 7 only depends on $(\boldsymbol{\alpha}, \mathbf{A})$ whereas the last two terms only depend on $(\boldsymbol{\mu}, \mathbf{B})$, and thus can be readily parallelized. Then, we find $(\boldsymbol{\mu}^*, \mathbf{B}^*)$ using spectral projected gradient descent [4], which works well in practice and achieves ε accuracy in $O(\log(1/\varepsilon))$ iterations, and find $(\boldsymbol{\alpha}^*, \mathbf{A}^*)$ analytically, since, for Gaussian sentiment distributions, the problem reduces to a least-square problem. Fortunately, in each subproblem, we can use the Markov property from Proposition 1 to precompute the sums and integrals in (8) in linear time, *i.e.*, $O(|\mathcal{H}_u(T)| + |\cup_{v \in \mathcal{N}(u)} \mathcal{H}_v(T)|)$. Appendix H summarizes the overall estimation algorithm.

2.2 Model simulation

We leverage the efficient sampling algorithm for multivariate Hawkes introduced by Farajtabar et al. [10] to design a scalable algorithm to sample opinions from our model. The two key ideas that allow us to adapt the procedure by Farajtabar et al. to our model of opinion dynamics, while keeping its efficiency, are as follows: (i) the opinion dynamics, defined by Eqs. 5 and 6, are Markovian and thus we can update individual intensities and opinions in $O(1)$ – let t_i and t_{i+1} be two consecutive events, then, we can compute $\lambda^*(t_{i+1})$ as $(\lambda^*(t_i) - \mu) \exp(-\nu(t_{i+1} - t_i)) + \mu$ and $x^*(t_{i+1})$ as $(x^*(t_i) - \alpha) \exp(-\omega(t_{i+1} - t_i)) + \alpha$, respectively; and, (ii) social networks are typically sparse and thus both \mathbf{A} and \mathbf{B} are also sparse, then, whenever a node expresses its opinion, only a small number of opinions and intensity functions in its local neighborhood will change. As a consequence, we can reuse the majority of samples from the intensity functions and sentiment distributions for the next new sample. Appendix I summarizes the overall simulation algorithm.

3 Opinion forecasting

Our goal here is developing efficient methods that leverage our model to forecast a user u 's opinion $x_u(t)$ at time t given the history $\mathcal{H}(t_0)$ up to time $t_0 < t$. In the context of our probabilistic model, we will forecast this opinion by efficiently computing the conditional expectation $\mathbb{E}_{\mathcal{H}(t) \setminus \mathcal{H}(t_0)}[x_u^*(t) | \mathcal{H}(t_0)]$, where $\mathcal{H}(t) \setminus \mathcal{H}(t_0)$ denotes the average across histories from t_0 to t , while conditioning on the history up to $\mathcal{H}(t_0)$.

To this aim, we will develop analytical and sampling based methods to compute the above conditional expectation. Moreover, we will use the former to identify under which conditions users' average opinion converges to a steady state and, if so, find the steady state opinion. In this section, we write $\mathcal{H}_t = \mathcal{H}(t)$ to lighten the notation and denote the eigenvalues of a matrix \mathbf{X} by $\xi(\mathbf{X})$.

3.1 Analytical forecasting

In this section, we derive a set of formulas to compute the conditional expectation for both Poisson and Hawkes messages intensities. However, since the derivation of such formulas for general multivariate Hawkes is difficult, we focus here on the case when $b_{vu} = 0$ for all $v, u \in \mathcal{G}, v \neq u$, and rely on the efficient sampling based method for the general case.

I. Poisson intensity. Consider each user's messages follow a Poisson process with rate μ_u . Then, the conditional average opinion is given by (proven in Appendix C):

Theorem 2. *Given a collection of messages \mathcal{H}_{t_0} recorded during a time period $[0, t_0)$ and $\lambda_u^*(t) = \mu_u$ for all $u \in \mathcal{G}$, then,*

$$\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}] = e^{(\mathbf{A}\boldsymbol{\Lambda}_1 - \omega\mathbf{I})(t-t_0)} \mathbf{x}(t_0) + \omega(\mathbf{A}\boldsymbol{\Lambda}_1 - \omega\mathbf{I})^{-1} \left(e^{(\mathbf{A}\boldsymbol{\Lambda}_1 - \omega\mathbf{I})(t-t_0)} - \mathbf{I} \right) \boldsymbol{\alpha}, \quad (9)$$

where $\mathbf{\Lambda}_1 := \text{diag}[\boldsymbol{\mu}]$ and $(\mathbf{x}(t_0))_{u \in \mathcal{V}} = \alpha_u + \sum_{v \in \mathcal{N}(u)} a_{uv} \sum_{t_i \in \mathcal{H}_v(t_0)} e^{-\omega(t_0-t_i)} m_v(t_i)$.

Remarkably, we can efficiently compute both terms in Eq. 9 by using the iterative algorithm by Al-Mohy et al. [2] for the matrix exponentials and the well-known GMRES method [22] for the matrix inversion. Given this predictive formula, we can easily study the stability condition and, for stable systems, find the steady state conditional average opinion (proven in Appendix D):

Theorem 3. *Given the conditions of Theorem 2, if $\text{Re}[\xi(\mathbf{A}\mathbf{\Lambda}_1)] < \omega$, then,*

$$\lim_{t \rightarrow \infty} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [\mathbf{x}^*(t) | \mathcal{H}_{t_0}] = \left(I - \frac{\mathbf{A}\mathbf{\Lambda}_1}{w} \right)^{-1} \boldsymbol{\alpha}. \quad (10)$$

The above results indicate that the conditional average opinions are nonlinearly related to the parameter matrix \mathbf{A} , which depends on the network structure, and the message rates $\boldsymbol{\mu}$, which in this case are assumed to be constant and independent on the network structure. Figure 1 provides empirical evidence of these results.

II. Multivariate Hawkes Process. Consider each user's messages follow a multivariate Hawkes process, given by Eq. 3, and $b_{vu} = 0$ for all $v, u \in \mathcal{G}, v \neq u$. Then, the conditional average opinion is given by (proven in Appendix E):

Theorem 4. *Given a collection of messages \mathcal{H}_{t_0} recorded during a time period $[0, t_0)$ and $\lambda_u^*(t) = \mu_u + b_{uu} \sum_{e_i \in \mathcal{H}_u(t)} e^{-\nu(t-t_i)}$ for all $u \in \mathcal{G}$, then, the conditional average satisfies the following differential equation:*

$$\frac{d\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [\mathbf{x}^*(t) | \mathcal{H}_{t_0}]}{dt} = [-\omega I + \mathbf{A}\mathbf{\Lambda}(t)] \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [\mathbf{x}^*(t) | \mathcal{H}_{t_0}] + \omega \boldsymbol{\alpha}, \quad (11)$$

where

$$\begin{aligned} \mathbf{\Lambda}(t) &= \text{diag} \left(\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [\boldsymbol{\lambda}^*(t) | \mathcal{H}_{t_0}] \right), \\ \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [\boldsymbol{\lambda}^*(t) | \mathcal{H}_{t_0}] &= e^{(\mathbf{B} - \nu I)(t-t_0)} \boldsymbol{\eta}(t_0) + \nu (\mathbf{B} - \nu I)^{-1} \left(e^{(\mathbf{B} - \nu I)(t-t_0)} - \mathbf{I} \right) \boldsymbol{\mu} \quad \forall t \geq t_0, \\ (\boldsymbol{\eta}(t_0))_{u \in \mathcal{V}} &= \mu_u + \sum_{v \in \mathcal{N}(u)} b_{uv} \sum_{t_i \in \mathcal{H}_v(t_0)} e^{-\nu(t_0-t_i)}, \\ \mathbf{B} &= \text{diag} \left([b_{11}, \dots, b_{|\mathcal{V}| |\mathcal{V}|}]^\top \right). \end{aligned}$$

Here, we can compute the conditional average by solving numerically the differential equation above, which is not stochastic, where we can efficiently compute the vector $\mathbb{E}_{\mathcal{H}_t} [\boldsymbol{\lambda}^*(t)]$ by using again the algorithm by Al-Mohy et al. [2] and the GMRES method [22].

In this case, the stability condition and the steady state conditional average opinion are given by (proven in Appendix F):

Theorem 5. *Given the conditions of Theorem 4, if the transition matrix $\Phi(t)$ associated to the time-varying linear system described by Eq. 11 satisfies that $\|\Phi(t)\| \leq \gamma e^{-ct} \forall t > 0$, where $\gamma, c > 0$, then,*

$$\lim_{t \rightarrow \infty} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [\mathbf{x}^*(t) | \mathcal{H}_{t_0}] = \left(I - \frac{\mathbf{A}\mathbf{\Lambda}_2}{w} \right)^{-1} \boldsymbol{\alpha}, \quad (12)$$

where $\mathbf{\Lambda}_2 := \text{diag} \left[\left(I - \frac{\mathbf{B}}{\nu} \right)^{-1} \boldsymbol{\mu} \right]$

The above results indicate that the conditional average opinions are nonlinearly related to the parameter matrices \mathbf{A} and \mathbf{B} . This suggests that the effect of the temporal influence on the opinion evolution, by means of the parameter matrix \mathbf{B} of the multivariate Hawkes process, is non trivial. We illustrate this result empirically in Figure 1.

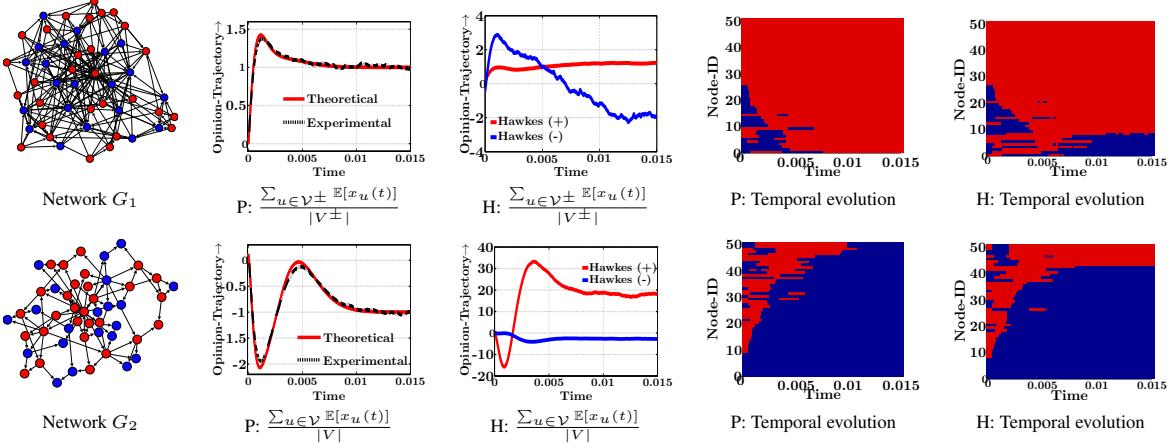


Figure 1: Opinion dynamics on two 50-node networks G_1 (top) and G_2 (bottom) for Poisson (P) and Hawkes (H) message intensities. The first column visualizes the two networks and opinion of each node at $t = 0$ (positive/negative opinions in red/blue). The second column shows the temporal evolution of the theoretical and empirical average opinion for Poisson intensities. The third column shows the temporal evolution of the empirical average opinion for Hawkes intensities, where we compute the average separately for positive (+) and negative (−) opinions in the steady state. The fourth and fifth columns shows the polarity of average opinion per user over time.

3.2 Simulation based forecasting

Given the efficient simulation procedure described in Section 2.2, we can readily derive a general simulation based formula for opinion forecasting:

$$\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [x^*(t) | \mathcal{H}_{t_0}] \approx \hat{x}^*(t) = \frac{1}{n} \sum_{l=1}^n x_l^*(t), \quad (13)$$

where n is the number of times that we simulate the opinion dynamics and $x_l^*(t)$ gathers the users' opinion at time t for the l -th simulation. Moreover, we have the following theoretical guarantee (proven in Appendix G):

Theorem 6. *Simulate the opinion dynamics up to time $t > t_0$ the following number of times:*

$$n \geq \frac{1}{3\epsilon^2} (6\sigma_{\max}^2 + 4x_{\max}\epsilon) \log(2/\delta), \quad (14)$$

where $\sigma_{\max}^2 = \max_{u \in \mathcal{G}} \sigma_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}^2(x_u^*(t) | \mathcal{H}_{t_0})$ is the maximum variance of the users' opinions, which we analyze in Appendix G, and $x_{\max} \geq |x_u(t)|, \forall u \in \mathcal{G}$ is an upper bound on the users' (absolute) opinions. Then, for each user $u \in \mathcal{G}$, the error between her true and estimated average opinion satisfies that $|\hat{x}_u^*(t) - \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [x_u^*(t) | \mathcal{H}_{t_0}]| \leq \epsilon$ with probability at least $1 - \delta$.

4 Experiments

4.1 Experiments on synthetic data

We first provide empirical evidence that our model is able to produce different types of opinion dynamics, which may or may not converge to a steady state of consensus or polarization. Then, we show that our model estimation and simulation algorithms as well as our predictive formulas scale to networks with millions of users and events. Appendix J contains an evaluation of the accuracy of our model parameter estimation method.

Different types of opinion dynamics. We first simulate our model on two different small networks using Poisson intensities, *i.e.*, $\lambda_u^*(t) = \mu_u$, $\mu_u \sim U(0, 1) \forall u$, and then simulate our model on the same networks while using Hawkes intensities with $b_{vu} \sim U(0, 1)$ on 5% of the nodes, chosen at random, and the original Poisson intensities on the remaining nodes. Figure 1 summarizes the results, which show that (i) our model is able to produce opinion

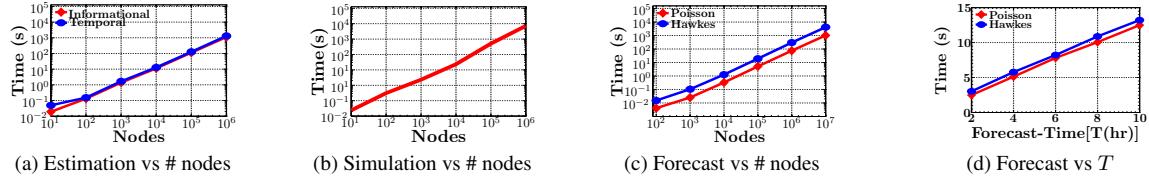


Figure 2: Panels (a) and (b) show running time of our estimation and simulation procedures against number of nodes, where the average number of events per node is 10. Panels (c) and (d) show the running time needed to compute our analytical formulas against number of nodes and time horizon $T = t - t_0$, where the number of nodes is 10^3 . In Panel (c), $T = 6$ hours. For all panels, the average degree per node is 30. The experiments are carried out in a single machine with 24 cores and 64 GB of main memory.

dynamics that converge to consensus (second column) and polarization (third column); (ii) the opinion forecasting formulas described in Section 3 closely match an simulation based estimation (second column); and, (iii) the evolution of the average opinion and whether opinions converge to a steady state of consensus or polarization depend on the functional form of message intensity⁵.

Scalability. Figure 2 shows that our model estimation and simulation algorithms, described in Sections 2.1 and 2.2, and our analytical predictive formulas, described in Section 3.1, scale to networks with millions of users and events. For example, our algorithm takes 20 minutes to estimate the model parameters from 10 million events generated by one million nodes using a single machine with 24 cores and 64 GB RAM.

4.2 Experiments on real data

We use real data gathered from Twitter to show that our model can forecast users’ opinions more accurately than six state of the art methods [6, 7, 8, 14, 18, 25].

Experimental Setup. We experimented with five Twitter datasets about current real-world events (Politics, Movie, Fight, Bollywood and US), in which, for each recorded message i , we compute its sentiment value m_i using a popular sentiment analysis toolbox, specially designed for Twitter [12]. Here, the sentiment takes values $m \in (-1, 1)$ and we consider the sentiment polarity to be simply $\text{sign}(m)$. Appendix K contains further details and statistics about these datasets.

Opinion forecasting. We first evaluate the performance of our model at predicting sentiment (expressed opinion) at a message level. To do so, for each dataset, we first estimate the parameters of our model, SLANT, using messages from a training set containing the (chronologically) first 90% of the messages. Here, we set the decay parameters of the exponential triggering kernels $\kappa(t)$ and $g(t)$ by cross-validation. Then, we evaluate the predictive performance of our opinion forecasting formulas using the last 10% of the messages⁶. More specifically, we predict the sentiment value m for each message posted by user u in the test set given the history up to T hours before the time of the message as $\hat{m} = E_{\mathcal{H}_t \setminus \mathcal{H}_{t-T}}[x_u^*(t) | \mathcal{H}_{t-T}]$. We compare the performance of our model with the asynchronous linear model (AsLM) [7], DeGroot’s model [8], the voter model [25], the biased voter model [6], the flocking model [14], and the sentiment prediction method based on collaborative filtering by Kim et al. [18], in terms of: (i) the mean squared error between the true (m) and the estimated (\hat{m}) sentiment value for all messages in the held-out set, *i.e.*, $\mathbb{E}[(m - \hat{m})^2]$, and (ii) the failure rate, defined as the probability that the true and the estimated polarity do not coincide, *i.e.*, $\mathbb{P}(\text{sign}(m) \neq \text{sign}(\hat{m}))$. For the baselines algorithms, which work in discrete time, we simulate N_T rounds in $(t - T, t)$, where N_T is the number of posts in time T . Figure 3 summarizes the results, which show that: (i) our opinion forecasting formulas consistently outperform others both in terms of MSE (often by an order of magnitude) and failure rate;⁷ (ii) its forecasting performance degrades gracefully with respect to T , in contrast, competing methods often fail catastrophically; and, (iii) it achieves an additional mileage by using Hawkes processes instead of Poisson processes. To some extent, we believe SLANT’s superior performance is due to its ability to leverage historical data to learn its model parameters and then simulate realistic temporal patterns.

⁵For these particular networks, Poisson intensities lead to consensus while Hawkes’ intensities lead to polarization, however, we did find other examples in which Poisson intensities lead to polarization and Hawkes intensities lead to consensus.

⁶Here, we do not distinguish between analytical and sampling based forecasting since, in practice, they closely match each other.

⁷The failure rate is very close to zero for those datasets in which most users post messages with the same polarity.

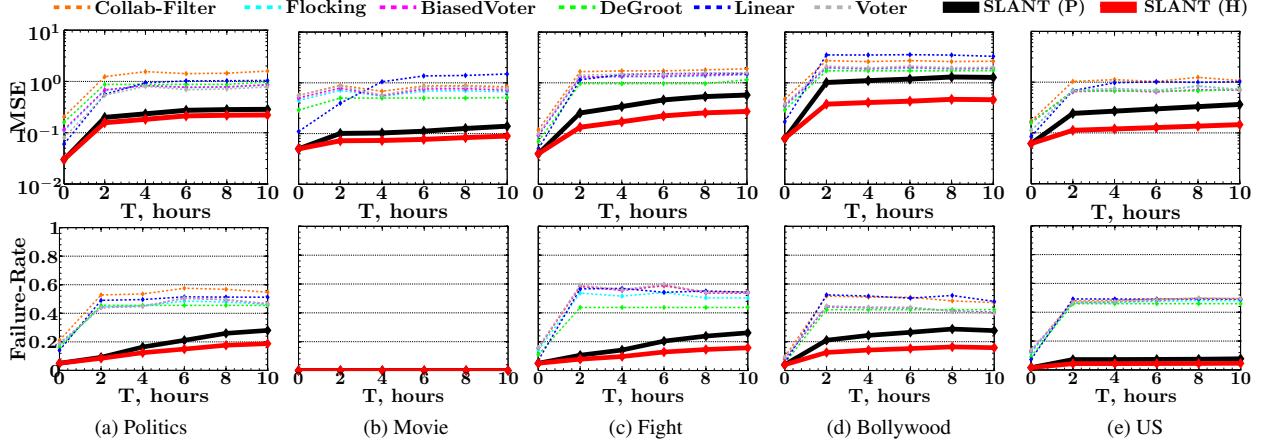


Figure 3: Sentiment prediction performance using a 10% held-out set for each real-world dataset. Performance is measured in terms of mean squared error (MSE) on the sentiment value, $\mathbb{E}[(m - \hat{m})^2]$, and failure rate on the sentiment polarity, $\mathbb{P}(\text{sign}(m) \neq \text{sign}(\hat{m}))$. For each message in the held-out set, we predict the sentiment value m given the history up to T hours before the time of the message, for different values of T . Nowcasting corresponds to $T = 0$ and forecasting to $T > 0$. The sentiment value $m \in (-1, 1)$ and the sentiment polarity $\text{sign}(m) \in \{-1, 1\}$.

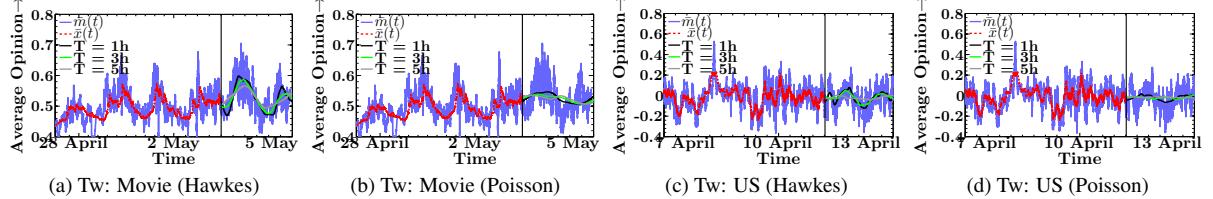


Figure 4: Macroscopic sentiment prediction given by our model for two real-world datasets. The panels show the observed sentiment $\bar{m}(t)$ (in blue, running average), inferred opinion $\bar{x}(t)$ on the training set (in red), and forecasted opinion $\mathbb{E}_{\mathcal{H}_{t-T}}[x_u(t)|\mathcal{H}_{t-T}]$ for $T = 1, 3$, and 5 hours on the test set (in black, green and gray, respectively), where the symbol $\bar{\cdot}$ denotes average across users.

Finally, we look at the forecasting results at a network level and show that our forecasting formulas can also predict the evolution of opinions macroscopically (in terms of the average opinion across users). Figure 4 summarizes the results for two real world datasets, which show that the forecasted opinions become less accurate as the time T becomes larger, since the average is computed on longer time periods. As expected, our model is more accurate when the message intensities are modeled using multivariate Hawkes. We found qualitatively similar results for the remaining datasets.

5 Conclusions

We proposed a modeling framework of opinion dynamics, whose key innovation is modeling users' *latent* opinions as continuous-time stochastic processes driven by a set of marked jump stochastic differential equations (SDEs) [13]. Such construction allows each user's latent opinion to be modulated over time by the opinions asynchronously *expressed* by her neighbors as *sentiment* messages. We then exploited a key property of our model, the Markov property, to design efficient parameter estimation and simulation algorithms, which scale to networks with millions of nodes. Moreover, we derived a set of novel predictive formulas for efficient and accurate opinion forecasting and identified conditions under which opinions converge to a steady state of consensus or polarization. Finally, we experimented with real data gathered from Twitter and showed that our framework achieves more accurate opinion forecasting than state-of-the-arts.

Our model opens up many interesting venues for future work. For example, in Eq. 4, our model assumes a linear dependence between users' opinions, however, in some scenarios, this may be a coarse approximation. A natural follow-up to improve the opinion forecasting accuracy would be considering nonlinear dependences between opinions.

It would be interesting to augment our model to jointly consider correlations between different topics. One could leverage our modeling framework to design opinion shaping algorithms based on stochastic optimal control [13, 24]. Finally, one of the key modeling ideas is realizing that users’ expressed opinions (be it in the form of thumbs up/down or text sentiment) can be viewed as noisy discrete samples of the users’ latent opinion localized in time. It would be very interesting to generalize this idea to any type of event data and derive sampling theorems and conditions under which an underlying general continuous signal of interest (be it user’s opinion, expertise, or wealth) can be recovered from event data with provable guarantees.

References

- [1] O. Aalen, Ø. Borgan, and H. Gjessing. *Survival and event history analysis: a process point of view*. Springer Verlag, 2008.
- [2] A. H. Al-Mohy and N. J. Higham. Computing the action of the matrix exponential, with an application to exponential integrators. *SIAM journal on scientific computing*, 33(2):488–511, 2011.
- [3] R. Axelrod. The dissemination of culture a model with local convergence and global polarization. *Journal of conflict resolution*, 41(2):203–226, 1997.
- [4] E. G. Birgin, J. M. Martínez, and M. Raydan. Nonmonotone spectral projected gradient methods on convex sets. *SIAM Journal on Optimization*, 10(4), 2000.
- [5] P. Clifford and A. Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973.
- [6] A. Das, S. Gollapudi, and K. Munagala. Modeling opinion dynamics in social networks. In *WSDM*, 2014.
- [7] A. De, S. Bhattacharya, P. Bhattacharya, N. Ganguly, and S. Chakrabarti. Learning a linear influence model from transient opinion dynamics. In *CIKM*, 2014.
- [8] M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345), 1974.
- [9] M. Farajtabar, N. Du, M. Gomez-Rodriguez, I. Valera, L. Song, and H. Zha. Shaping social activity by incentivizing users. In *NIPS*, 2014.
- [10] M. Farajtabar, Y. Wang, M. Gomez-Rodriguez, S. Li, H. Zha, and L. Song. Coevolve: A joint point process model for information diffusion and network co-evolution. In *NIPS*, 2015.
- [11] M. Gomez-Rodriguez, D. Balduzzi, and B. Schölkopf. Uncovering the Temporal Dynamics of Diffusion Networks. In *ICML*, 2011.
- [12] A. Hannak, E. Anderson, L. F. Barrett, S. Lehmann, A. Mislove, and M. Riedewald. Tweetin’ in the rain: Exploring societal-scale effects of weather on mood. In *ICWSM*, 2012.
- [13] F. B. Hanson. *Applied Stochastic Processes and Control for Jump-Diffusions*. SIAM, 2007.
- [14] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3), 2002.
- [15] D. Hinrichsen, A. Ilchmann, and A. Pritchard. Robustness of stability of time-varying linear systems. *Journal of Differential Equations*, 82(2):219 – 250, 1989.
- [16] P. Holme and M. E. Newman. Nonequilibrium phase transition in the coevolution of networks and opinions. *Physical Review E*, 74(5):056108, 2006.
- [17] T. Karppi and K. Crawford. Social media, financial algorithms and the hack crash. *TC&S*, 2015.
- [18] J. Kim, J.-B. Yoo, H. Lim, H. Qiu, Z. Kozareva, and A. Galstyan. Sentiment prediction using collaborative filtering. In *ICWSM*, 2013.

- [19] J. Leskovec, D. Chakrabarti, J. M. Kleinberg, C. Faloutsos, and Z. Ghahramani. Kronecker graphs: An approach to modeling networks. *JMLR*, 2010.
- [20] B. Pang and L. Lee. Opinion mining and sentiment analysis. *F&T in information retrieval*, 2(1-2), 2008.
- [21] B. H. Raven. The bases of power: Origins and recent developments. *Journal of social issues*, 49(4), 1993.
- [22] Y. Saad and M. H. Schultz. Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on scientific and statistical computing*, 7(3):856–869, 1986.
- [23] I. Valera and M. Gomez-Rodriguez. Modeling adoption and usage of competing products. In *Proceedings of the 2015 IEEE International Conference on Data Mining*, 2015.
- [24] Y. Wang, E. Theodorou, A. Verma, and L. Song. Steering opinion dynamics in information diffusion networks. *arXiv preprint arXiv:1603.09021*, 2016.
- [25] M. E. Yildiz, R. Pagliari, A. Ozdaglar, and A. Scaglione. Voting models in random networks. In *Information Theory and Applications Workshop*, pages 1–7, 2010.

A Proof of Proposition 1

Given $\mathbf{x}^*(t) := [\mathbf{x}_1(t), \dots, \mathbf{x}_{|\mathcal{V}|}(t)]^T$ and $\boldsymbol{\lambda}^*(t) := [\boldsymbol{\lambda}_1(t), \dots, \boldsymbol{\lambda}_{|\mathcal{V}|}(t)]^T$, we can compactly rewrite Eqs. 3 and 4 for all users as:

$$\mathbf{x}^*(t) = \boldsymbol{\alpha} + \mathbf{A} \int_0^t g(t-\theta) \mathbf{m}(t) \odot d\mathbf{N}(t) \quad (15)$$

and

$$\boldsymbol{\lambda}^*(t) = \boldsymbol{\mu} + \mathbf{B} \int_0^t \kappa(t-\theta) \odot d\mathbf{N}(t). \quad (16)$$

Then, it easily follows that

$$d\mathbf{x}^*(t) = \mathbf{A} \int_0^t g'(t-\theta) \mathbf{m}(t) \odot d\mathbf{N}(t) + g(0) \mathbf{A} \mathbf{m}(t) \odot d\mathbf{N}(t), \quad (17)$$

where $g(t) = e^{-\omega t}$ and $g'(t-\theta) = -\omega g(t-\theta)$. And, we can rewrite the above expression as

$$d\mathbf{x}^*(t) = \omega(\boldsymbol{\alpha} - \mathbf{x}^*(t))dt + \mathbf{A}(\mathbf{m}(t) \odot d\mathbf{N}(t)). \quad (18)$$

Similarly, we can show that

$$d\boldsymbol{\lambda}^*(t) = \nu(\boldsymbol{\mu} - \boldsymbol{\lambda}^*(t))dt + \mathbf{B} d\mathbf{N}(t).$$

B Auxiliary theoretical results

The proofs of Theorems 2 and 5 rely on the following auxiliary Lemma.

Lemma 7. *The expected opinion $\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}]$ in the model of opinion dynamics defined by Eqs. 4 and 3 with exponential triggering kernels with parameters ω and ν satisfies the following differential equation:*

$$\frac{d\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}]}{dt} = \mathbf{A} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) \odot \boldsymbol{\lambda}^*(t) | \mathcal{H}_{t_0}] - \omega \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}] + \omega \boldsymbol{\alpha}, \quad (19)$$

where $\mathbf{A} = (a_{vu})_{v,u \in \mathcal{G}}$ and the sign \odot denotes pointwise product.

Using that $\mathbb{E}[m_v(\theta) | x_v^*(\theta)] = x_v^*(\theta)$, we can compute the average opinion of user u across all possible histories from Eq. 4 as

$$\begin{aligned} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[x_u^*(t) | \mathcal{H}_{t_0}] &= \alpha_u + \sum_{v \in \mathcal{N}(u)} a_{uv} \int_0^t g(t-\theta) \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[m_v(\theta) dN_v(\theta) | \mathcal{H}_{t_0}] \\ &= \alpha_u + \sum_{v \in \mathcal{N}(u)} a_{uv} \sum_{t_i \in \mathcal{H}_v(t_0)} g(t-t_i) m_v(t_i) + \sum_{v \in \mathcal{N}(u)} a_{uv} \int_{t_0}^t g(t-\theta) \mathbb{E}_{\mathcal{H}(\theta) \setminus \mathcal{H}_{t_0}}[x_v^*(\theta) \lambda_v^*(\theta) | \mathcal{H}_{t_0}] d\theta, \end{aligned}$$

and we can write the expectation of the opinion for all users in vectorial form as

$$\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t)] = \mathbf{v}(t) + \mathbf{A} \int_0^t g(t-\theta) \mathbb{E}_{\mathcal{H}(\theta) \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(\theta) \odot \boldsymbol{\lambda}^*(\theta)] d\theta, \quad (20)$$

where the \odot denotes pointwise product and

$$(\mathbf{v}(t))_u = \alpha_u + \sum_{v \in \mathcal{N}(u)} a_{uv} \sum_{t_i \in \mathcal{H}_v(t_0)} g(t-t_i) m_v(t_i).$$

Since $g(t) = e^{-\omega t}$, one may observe that $\omega \mathbf{v}(t) + \dot{\mathbf{v}}(t) = \omega \boldsymbol{\alpha}$. Then, by differentiating Eq. 20, we obtain

$$\frac{d\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}]}{dt} = \mathbf{A} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) \odot \boldsymbol{\lambda}^*(t) | \mathcal{H}_{t_0}] - \omega \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}] + \omega \boldsymbol{\alpha}, \quad (21)$$

C Proof of Theorem 2

Using Lemma 7 and $\lambda_u^*(t) = \mu_u$, we obtain

$$\frac{d\mathbb{E}_{\mathcal{H}_t}[\mathbf{x}^*(t)]}{dt} = [-\omega I + \mathbf{A}\mathbf{\Lambda}_1]\mathbb{E}_{\mathcal{H}_t}[\mathbf{x}^*(t)] + \omega\boldsymbol{\alpha}, \quad (22)$$

where $\mathbf{\Lambda}_1 = \text{diag}[\boldsymbol{\mu}]$. Then, we apply the Laplace transform to the expression above and obtain

$$\hat{\mathbf{x}}(s) = [sI + \omega I - A\mathbf{\Lambda}_1]^{-1}\mathbf{x}(t_0) + \frac{\omega}{s}[sI + \omega I - A\mathbf{\Lambda}_1]^{-1}\boldsymbol{\alpha},$$

where we leverage the fact that, conditioning the prior history, the opinion is non-random, *i.e.*, $\mathbb{E}_{\mathcal{H}_{t_0} \setminus \mathcal{H}_{t_0^-}}[\mathbf{x}(t_0) | \mathcal{H}_{t_0^-}] = \mathbf{x}(t_0)$. Finally, applying the inverse Laplace transform, we obtain the average opinion $\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}]$ in time domain as

$$\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}] = e^{(\mathbf{A}\mathbf{\Lambda}_1 - \omega I)(t-t_0)}\mathbf{x}(t_0) + \omega(\mathbf{A}\mathbf{\Lambda}_1 - \omega I)^{-1}\left(e^{(\mathbf{A}\mathbf{\Lambda}_1 - \omega I)(t-t_0)} - \mathbf{I}\right)\boldsymbol{\alpha}.$$

D Proof of Theorem 3

Theorem 2 states that the average users' opinion $\mathbb{E}_{\mathcal{H}_t}[\mathbf{x}^*(t)]$ in time domain is given by

$$\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}] = e^{(\mathbf{A}\mathbf{\Lambda}_1 - \omega I)(t-t_0)}\mathbf{x}(t_0) + \omega(\mathbf{A}\mathbf{\Lambda}_1 - \omega I)^{-1}\left(e^{(\mathbf{A}\mathbf{\Lambda}_1 - \omega I)(t-t_0)} - \mathbf{I}\right)\boldsymbol{\alpha}.$$

If $\text{Re}[\xi(\mathbf{A}\mathbf{\Lambda}_1)] < \omega$, where $\xi(\mathbf{X})$ denote the eigenvalues of matrix \mathbf{X} , it easily follows that

$$\lim_{t \rightarrow \infty} \mathbb{E}_{\mathcal{H}_t}[\mathbf{x}^*(t)] = \left(I - \frac{\mathbf{A}\mathbf{\Lambda}_1}{\omega}\right)^{-1}\boldsymbol{\alpha}. \quad (23)$$

E Proof of Theorem 4

Assume $b_{vu} = 0$ for all $v, u \in \mathcal{V}, v \neq u$. Then, $\lambda_v^*(t)$ only depends on user v 's history and, since $x_v^*(t)$ only depends on the history of the user v 's neighbors $\mathcal{N}(v)$, we can write

$$\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) \odot \boldsymbol{\lambda}^*(t) | \mathcal{H}_{t_0}] = \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}] \odot \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\boldsymbol{\lambda}^*(t)],$$

and rewrite Eq. 21 as

$$\begin{aligned} \frac{d\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}]}{dt} = \\ \mathbf{A}(\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}] \odot \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\boldsymbol{\lambda}^*(t) | \mathcal{H}_{t_0}]) - \omega\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}] + \omega\boldsymbol{\alpha}. \end{aligned} \quad (24)$$

We can now compute $\mathbb{E}_{\mathcal{H}(\theta) \setminus \mathcal{H}_{t_0}}[\boldsymbol{\lambda}^*(\theta) | \mathcal{H}_{t_0}]$ analytically as follows. From Eq. 3, we obtain

$$\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\boldsymbol{\lambda}^*(t) | \mathcal{H}_{t_0}] = \boldsymbol{\eta}(t) + \int_{t_0}^t \mathbf{B}\kappa(t-\theta)\mathbb{E}_{\mathcal{H}_\theta \setminus \mathcal{H}_{t_0}}[\boldsymbol{\lambda}^*(\theta)]d\theta, \quad (25)$$

where $\kappa(t) = e^{-\nu t}$, $[\boldsymbol{\eta}(t)]_{u \in \mathcal{V}} = \mu_u + \sum_{v \in \mathcal{N}(u)} b_{uv} \sum_{i \in \mathcal{H}_v(t_0)} \kappa(t - t_i)$ and $\mathbf{B} = (b_{vu})_{v, u \in \mathcal{V}}$, where $b_{vu} = 0$ for all $v \neq u$, by assumption. Differentiating with respect to t , we get,

$$\frac{d}{dt}\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\boldsymbol{\lambda}^*(t) | \mathcal{H}_{t_0}] = -\nu\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\boldsymbol{\lambda}^*(t) | \mathcal{H}_{t_0}] + \mathbf{B}\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\boldsymbol{\lambda}^*(t) | \mathcal{H}_{t_0}] + \nu\boldsymbol{\mu},$$

with initial condition $\mathbb{E}_{\mathcal{H}(t_0^+) \setminus \mathcal{H}(t_0)}[\boldsymbol{\lambda}^*(t_0)] = \boldsymbol{\eta}(t_0)$. By taking the Laplace transform and then applying inverse Laplace transform,

$$\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\boldsymbol{\lambda}^*(t)|\mathcal{H}_{t_0}] = e^{(\mathbf{B} - \nu I)(t-t_0)} \boldsymbol{\eta}(t_0) + \nu(\mathbf{B} - \nu I)^{-1} \left(e^{(\mathbf{B} - \nu I)(t-t_0)} - \mathbf{I} \right) \boldsymbol{\mu} \quad \forall t \geq t_0, \quad (26)$$

where $\boldsymbol{\eta}(t_0) = \mathbb{E}_{\mathcal{H}(t_0^+) \setminus \mathcal{H}(t_0)}[\boldsymbol{\lambda}^*(t_0)]$. Using Eqs. 24 and 26, as well as $\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t)] \odot \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\boldsymbol{\lambda}^*(t)] = \boldsymbol{\Lambda}(t) \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t)]$, where $\boldsymbol{\Lambda}(t) := \text{diag}[\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\boldsymbol{\lambda}^*(t)]]$, we obtain

$$\frac{d\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t)]}{dt} = [-\omega I + \mathbf{A}\boldsymbol{\Lambda}(t)]\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t)] + \omega\boldsymbol{\alpha},$$

with initial conditions $(\mathbb{E}_{\mathcal{H}(t_0^+) \setminus \mathcal{H}(t_0)}[\mathbf{x}^*(t_0)])_{u \in \mathcal{V}} = \alpha_u + \sum_{v \in \mathcal{N}(u)} a_{uv} \sum_{t_i \in \mathcal{H}_v(t_0)} g(t_0 - t_i) m_v(t_i)$.

F Proof of Theorem 5

Theorem 4 states that the average users' opinion $\mathbb{E}_{\mathcal{H}_t}[\mathbf{x}^*(t)]$ in time domain is given by

$$\frac{d\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t)|\mathcal{H}_{t_0}]}{dt} = [-\omega I + \mathbf{A}\boldsymbol{\Lambda}(t)]\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t)|\mathcal{H}_{t_0}] + \omega\boldsymbol{\alpha}. \quad (27)$$

In such systems, solutions can be written as [15]

$$\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t)|\mathcal{H}_{t_0}] = \Phi(t)\boldsymbol{\alpha} + \omega \int_0^t \Phi(s)\boldsymbol{\alpha} ds, \quad (28)$$

where the transition matrix $\Phi(t)$ defines as a solution of the matrix differential equation

$$\dot{\Phi}(t) = [-\omega I + \mathbf{A}\boldsymbol{\Lambda}(t)]\Phi(t) \text{ with } \Phi(0) = \mathbf{I}.$$

If $\Phi(t)$ satisfies $\|\Phi(t)\| \leq \gamma e^{-ct} \forall t > 0$ for $\gamma, c > 0$ then the steady state solution to Eq. 28 is given by [15]

$$\lim_{t \rightarrow \infty} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t)|\mathcal{H}_{t_0}] = \left(\mathbf{I} - \frac{\mathbf{A}\boldsymbol{\Lambda}_2}{\omega} \right)^{-1} \boldsymbol{\alpha}.$$

where $\boldsymbol{\Lambda}_2 = \lim_{t \rightarrow \infty} \boldsymbol{\Lambda}(t) = \text{diag} \left[\mathbf{I} - \frac{\mathbf{B}}{\nu} \right]^{-1} \boldsymbol{\mu}$.

G Proof of Theorem 6

Let $\{\mathbf{x}_l^*(t)\}_{l=1}^n$ be the simulated opinions for all users and define $\mathbf{s}(t) = \frac{1}{n} \sum_{l=1}^n \mathbf{s}_l(t)$, where $\mathbf{s}_l(t) = (\hat{x}_l^*(t) - \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}_l^*(t)|\mathcal{H}_{t_0}])$. Clearly, for a given t , all elements in $\mathbf{s}_l(t)$ are i.i.d. random variables with zero mean and variance, and we can bound $|s_u(t)| < 2x_{\max}$. Then, by Bernstein's inequality, the following holds true,

$$\mathbb{P}(|s_u(t)| > \epsilon) = \mathbb{P}(s_u(t) > \epsilon) + \mathbb{P}(s_u(t) < -\epsilon) > 2 \cdot \exp \left(- \frac{3n\epsilon^2}{6\sigma_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}^2(x_u^*(t)|\mathcal{H}_{t_0}) + 4x_{\max}\epsilon} \right)$$

Let $\sigma_{\max}^2(t) = \max_{u \in \mathcal{G}} \sigma_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}^2(x_u^*(t)|\mathcal{H}_{t_0})$. If we choose,

$$\delta < 2 \cdot \exp \left(- \frac{3n\epsilon^2}{6\sigma_{\max}^2(t) + 4x_{\max}\epsilon} \right) \quad (29)$$

we obtain the required bound for n . Moreover, given this choice of δ , we have $\mathbb{P}(|s_u(t)| < \epsilon) > 1 - \delta$ immediately.

However, a finite bound on n requires the variance $\sigma_{\max}^2(t)$ to be bounded for all t . Hence, we analyze the variance and its stability below.

G.1 Dynamics of variance

In this section we compute the time-domain evolution of the variance and characterize its stability for Poisson driven opinion dynamics. A general analysis of the variance for multidimensional Hawkes is left for future work.

Lemma 8. *Given a collection of messages \mathcal{H}_{t_0} recorded during a time period $[0, t_0)$ and $\lambda_u^*(t) = \mu_u$ for all $u \in \mathcal{G}$, the covariance matrix $\mathbf{\Gamma}(t_0, t)$ at any time t conditioned on the history \mathcal{H}_{t_0} can be described as,*

$$\text{vec}(\mathbf{\Gamma}(t_0, t)) = \int_0^t \Phi(t-\theta) \text{vec}[\sigma^2 \mathbf{A} \mathbf{\Lambda} \mathbf{A}^T + \mathbf{A} \text{diag}(\mathbb{E}_{\mathcal{H}_{t-\theta}}[\mathbf{x}^*(\theta)])^2 \mathbf{\Lambda} \mathbf{A}^T] d\theta.$$

where

$$\Phi(t) = e^{((-\omega I + \mathbf{A} \mathbf{\Lambda}) \otimes I + I \otimes (-\omega I + \mathbf{A} \mathbf{\Lambda}) + (\mathbf{A} \otimes \mathbf{A}) \widehat{\mathbf{\Lambda}})t},$$

$\widehat{\mathbf{\Lambda}}_{i^2, i^2} = \mathbf{\Lambda}^*(i)$, and $\mathbf{\Lambda} := \text{diag}[\mathbf{\lambda}]$. Moreover, the stability of the system is characterized by

$$\xi [(-\omega I + \mathbf{A} \mathbf{\Lambda}) \otimes I + I \otimes (-\omega I + \mathbf{A} \mathbf{\Lambda}) + (\mathbf{A} \otimes \mathbf{A}) \widehat{\mathbf{\Lambda}}] < 0.$$

Proof. By definition, the covariance matrix is given by

$$\mathbf{\Gamma}(t_0, t) := \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[(\mathbf{x}^*(t) - \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t))) (\mathbf{x}^*(t) - \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t)))^T | \mathcal{H}_{t_0}]. \quad (30)$$

Hence, if we define $\Delta \mathbf{x} = (\mathbf{x}^*(t) - \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t)))$, we can compute the differential of the covariance matrix as

$$d\mathbf{\Gamma}(t_0, t) = \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[d(\Delta \mathbf{x} \Delta \mathbf{x}^T) | \mathcal{H}_{t_0}] = \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\Delta \mathbf{x} d(\Delta \mathbf{x}^T) + d(\Delta \mathbf{x}) \Delta \mathbf{x}^T + d(\Delta \mathbf{x}) d(\Delta \mathbf{x}^T) | \mathcal{H}_{t_0}], \quad (31)$$

where

$$d(\Delta \mathbf{x}) = d(\mathbf{x}^*(t) - \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t))) = d(\mathbf{x}^*(t)) - d(\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t))). \quad (32)$$

Next, note that

$$\mathbb{E}(d\mathbf{N}(t)d\mathbf{N}^T(t)) = \mathbb{E}[dN_i(t)dN_j(t)]_{i,j \in \mathcal{V}} = \mathbb{E}[\text{diag}(d\mathbf{N}(t))] = \mathbf{\Lambda}, \quad (33)$$

where the off-diagonal entries vanish, since two jumps cannot happen at the same time point [13]. Now, recall the Markov representation of our model, *i.e.*,

$$d\mathbf{x}^*(t) = -\omega \mathbf{x}^*(t)dt + \mathbf{A} \mathbf{M}^*(t)d\mathbf{N}(t) + \boldsymbol{\alpha} dt, \quad (34)$$

where $\mathbf{M}^*(t) := \text{diag}[\mathbf{m}(t)]$ is the diagonal formed by the sentiment vector and note that,

$$\mathbf{m}(t) \odot d\mathbf{N}(t) = \mathbf{M}^*(t)d\mathbf{N}(t) = \text{diag}[d\mathbf{N}(t)]\mathbf{m}(t), \quad (35)$$

and, using Eq. 21,

$$d\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}] = -\omega \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}]dt + \mathbf{A} \mathbf{\Lambda} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}]dt + \omega \boldsymbol{\alpha} dt. \quad (36)$$

Then, if we substitute Eqs. 34 and 36 in Eq. 32, we obtain

$$d(\Delta \mathbf{x}) = -\omega [\mathbf{x}^*(t) - \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t) | \mathcal{H}_{t_0})]dt + \mathbf{A}[\mathbf{M}^*(t)d\mathbf{N}(t) - \mathbf{\Lambda} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t) | \mathcal{H}_{t_0})dt]. \quad (37)$$

As a result, we can write

$$\begin{aligned} d\mathbf{\Gamma}(t_0, t) &= \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} \left[-2\omega \Delta \mathbf{x} \Delta \mathbf{x}^T dt + \underbrace{\Delta \mathbf{x} ((\mathbf{M}^*(t)d\mathbf{N}(t) - \mathbf{\Lambda} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t) | \mathcal{H}_{t_0})dt)^T \mathbf{A}^T)}_{\text{Term 1}} \right. \\ &\quad + \mathbf{A}(\mathbf{M}^*(t)d\mathbf{N}(t) - \mathbf{\Lambda} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t) | \mathcal{H}_{t_0})dt) \Delta \mathbf{x}^T \\ &\quad \left. + \underbrace{\mathbf{A}(\mathbf{M}^*(t)d\mathbf{N}(t) - \mathbf{\Lambda} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t) | \mathcal{H}_{t_0})dt)(\mathbf{M}^*(t)d\mathbf{N}(t) - \mathbf{\Lambda} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t) | \mathcal{H}_{t_0})dt)^T \mathbf{A}^T}_{\text{Term 2}} | \mathcal{H}_{t_0} \right], \end{aligned} \quad (38)$$

where Term 1 gives

$$\begin{aligned}
& \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [\Delta x(\mathbf{M}^*(t)d\mathbf{N}(t) - \mathbf{\Lambda} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t))dt)^T \mathbf{A}^T | \mathcal{H}_{t_0}] \\
&= \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [\Delta x(m^*(t) - \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t))dt)^T \mathbf{A}^T | \mathcal{H}_{t_0}] \mathbf{\Lambda} dt \quad (\text{Using Eq. 35 and the fact that } \mathbb{E}[\text{diag}(d\mathbf{N}(t))] = \mathbf{\Lambda}) \\
&= \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [\Delta x \cdot \mathbb{E}[(m^*(t) - \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t)))^T \mathbf{A}^T | x(t), \mathcal{H}_{t_0}]] \mathbf{\Lambda} dt \\
&= \mathbf{\Gamma}(t_0, t) \mathbf{\Lambda} dt,
\end{aligned}$$

and Term 2 gives

$$\begin{aligned}
& \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} \mathbf{A} [(\mathbf{M}^*(t)d\mathbf{N}(t) - \mathbf{\Lambda} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t))dt)(\mathbf{M}^*(t)d\mathbf{N}(t) - \mathbf{\Lambda} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t))dt)^T | \mathcal{H}_{t_0}] \mathbf{A}^T \\
&= \mathbf{A} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [\mathbf{M}^*(t) \mathbb{E}(d\mathbf{N}(t)d\mathbf{N}^T(t)) \mathbf{M}^*(t) | \mathcal{H}_{t_0}] \mathbf{A}^T + O(dt^2) \\
&= \mathbf{A} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}} [\mathbf{M}^2(t) | \mathcal{H}_{t_0}] \mathbf{\Lambda} dt \mathbf{A}^T \quad (\text{From Eq. 33}) \\
&= \mathbf{A} (\sigma^2 I + \text{diag}(\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(x(t)^2) | \mathcal{H}_{t_0})) \mathbf{\Lambda} \mathbf{A}^T dt \\
&= \mathbf{A} [\sigma^2 I + \mathbf{\Gamma}_{ii}(t_0, t) + \text{diag}[\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t) | \mathcal{H}_{t_0})]^2] \mathbf{\Lambda} \mathbf{A}^T dt.
\end{aligned}$$

Hence, substituting the expectations above into Eq. 38, we obtain

$$\begin{aligned}
\frac{d\mathbf{\Gamma}(t_0, t)}{dt} &= -2\omega \mathbf{\Gamma}(t_0, t) + \mathbf{\Gamma}(t_0, t) \mathbf{\Lambda} \mathbf{A}^T + \mathbf{A} \mathbf{\Lambda} \mathbf{\Gamma}(t_0, t) \\
&+ \mathbf{A} \mathbf{\Gamma}_{ii}(t) \mathbf{\Lambda} \mathbf{A}^T + \sigma^2 \mathbf{A} \mathbf{\Lambda} \mathbf{A}^T + \mathbf{A} \text{diag}(\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}])^2 \mathbf{\Lambda} \mathbf{A}^T,
\end{aligned} \tag{39}$$

Eq. 39 can be readily written in vectorial form by exploiting the properties of the Kronecker product as

$$\frac{d[\text{vec}(\mathbf{\Gamma}(t_0, t))]}{dt} = \mathbf{V}(t) \text{vec}(\mathbf{\Gamma}(t_0, t)) + \text{vec}(\sigma^2 \mathbf{A} \mathbf{\Lambda} \mathbf{A}^T + \mathbf{A} \text{diag}(\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(t) | \mathcal{H}_{t_0}])^2 \mathbf{\Lambda} \mathbf{A}^T)$$

where

$$\mathbf{V}(t) = (-\omega I + \mathbf{A} \mathbf{\Lambda}) \otimes I + I \otimes (-\omega I + \mathbf{A} \mathbf{\Lambda}) + \sum_{i=1}^n (\mathbf{A} \mathbf{\Lambda} P_i \otimes \mathbf{A} P_i),$$

$P_i = [\delta_{ii}]$ and \otimes stands for the Kronecker product. Hence, the closed form solution of the above equation can be written as,

$$\text{vec}(\mathbf{\Gamma}(t_0, t)) = \int_0^t e^{\mathbf{V}(t-\theta)} \text{vec}[\sigma^2 \mathbf{A} \mathbf{\Lambda} \mathbf{A}^T + \mathbf{A} \text{diag}(\mathbb{E}_{\mathcal{H}(\theta) \setminus \mathcal{H}_{t_0}}[\mathbf{x}^*(\theta) | \mathcal{H}_{t_0}])^2 \mathbf{\Lambda} \mathbf{A}^T] d\theta. \tag{40}$$

Moreover, the covariance matrix $\mathbf{\Gamma}(t_0, t)$ is bounded iff

$$\text{Re}[\lambda[(-\omega I + \mathbf{A} \mathbf{\Lambda}) \otimes I + I \otimes (-\omega I + \mathbf{A} \mathbf{\Lambda}) + (\mathbf{A} \otimes \mathbf{A}) \widehat{\mathbf{\Lambda}}]] < 0,$$

where $\widehat{\mathbf{\Lambda}} := \sum_{i=1}^{|\mathcal{V}|} \mathbf{\Lambda} P_i \otimes P_i$.

In this case, the steady state solution $\mathbf{\Gamma} = \lim_{t \rightarrow \infty} \mathbf{\Gamma}(t_0, t)$ is given by

$$-2\omega \mathbf{\Gamma} + \mathbf{\Gamma} \mathbf{\Lambda} \mathbf{A}^T + \mathbf{A} \mathbf{\Lambda} \mathbf{\Gamma} + \mathbf{A} \text{diag}(\mathbf{\Gamma}) \mathbf{\Lambda} \mathbf{A}^T + \sigma^2 \mathbf{A} \mathbf{\Lambda} \mathbf{A} + \mathbf{A} \text{diag}(\mathbb{E}_{\mathcal{H}_{t-}}[\mathbf{x}^*])^2 \mathbf{\Lambda} \mathbf{A}^T = 0. \tag{41}$$

that means $\lim_{t \rightarrow \infty} \text{vec}[\mathbf{\Gamma}(t_0, t)]$ is same as,

$$\begin{aligned}
\text{vec}(\mathbf{\Gamma}) &= [(-\omega I + \mathbf{A} \mathbf{\Lambda}) \otimes I + I \otimes (-\omega I + \mathbf{A} \mathbf{\Lambda})^T + (\mathbf{A} \otimes \mathbf{A}) \widehat{\mathbf{\Lambda}}]^{-1} \\
&\quad \times \text{vec}[\sigma^2 \mathbf{A} \mathbf{\Lambda} \mathbf{A}^T + \mathbf{A} \text{diag}(\mathbf{x}_\infty^*)^2 \mathbf{\Lambda} \mathbf{A}^T]
\end{aligned} \tag{42}$$

where $\mathbf{x}^* = \lim_{t \rightarrow \infty} \mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}(\mathbf{x}^*(t) | \mathcal{H}_{t_0})$. Finally, the variance of node u 's opinion at any time t is given by

$$\sigma_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}^2(x_u^*(t) | \mathcal{H}_{t_0}) = [\mathbf{\Gamma}(t_0, t)]_{u,u} = \text{vec}(P_u)^T \text{vec}(\mathbf{\Gamma}(t_0, t)),$$

where P_u is the sparse matrix with its only (u, u) entry to be equal to unity.

H Parameter estimation algorithm

H.1 Estimation of α and A

Algorithm H.1 summarizes the estimation of the opinion dynamics, *i.e.*, α and A , which reduces to a least-square problem.

Algorithm 1 Estimation of α and A

```

1: Input:  $\mathcal{H}_T, G(V, E)$ , regularizer  $\lambda$ , error bound  $\epsilon$ 
2: Output:  $\hat{\alpha}, \hat{A}$ 
3: Initialize:
4: IndexForV[1 : |V|] =  $\vec{0}$ 
5: for  $i = 1$  to  $|\mathcal{H}(T)|$  do
6:    $\mathcal{T}[u_i](\text{IndexForV}[u_i]) = (t_i, u_i, m_i)$ 
7:   IndexForV[ $u_i$ ]++
8: for  $u \in \mathcal{V}$  do
9:    $i = 0$ 
10:   $S = \mathcal{T}[u]$ 
11:  for  $v \in \mathcal{N}(u)$  do
12:     $S = \text{MergeTimes}(S, \mathcal{T}[v])$ 
13:    // Merges two sets w.r.t  $t_i$ , takes  $O(|\mathcal{H}_u(T)| + |\cup_{v \in \mathcal{N}(u)} \mathcal{H}_v(T)|)$  steps
14:     $\mathbf{x}_{\text{last}} = 0$ 
15:     $t_{\text{last}} = 0$ 
16:     $j = 0$ 
17:    for  $i = 1$  to  $|S|$  do
18:      // This loop takes  $O(|S|) = O(|\mathcal{H}_u(T)| + |\cup_{v \in \mathcal{N}(u)} \mathcal{H}_v(T)|)$  steps
19:       $(t_v, v, m_v) = S[i]$ 
20:       $t_{\text{now}} = t_v$ 
21:      if  $u = v$  then
22:         $\mathbf{x}_{\text{now}} = \mathbf{x}_{\text{last}} e^{-\omega(t_{\text{now}} - t_{\text{last}})}$ 
23:         $j++$ 
24:         $g[u](j, v) = \mathbf{x}_{\text{now}}$ 
25:         $Y[u](j) = m_u$ 
26:      else
27:         $\mathbf{x}_{\text{now}} = \mathbf{x}_{\text{last}} e^{-\omega(t_{\text{now}} - t_{\text{last}})} + m_v$ 
28:         $t_{\text{last}} = t_{\text{now}}$ 
29:         $\mathbf{x}_{\text{last}} = \mathbf{x}_{\text{now}}$ 
30: // Estimation of  $(\alpha, A)$ 
31: for  $u \in \mathcal{V}$  do
32:    $a = \text{InferOpinionParams}(Y[u], \lambda, g[u])$ 
33:    $\hat{\alpha}_u = a[1]$ 
34:    $\hat{A}[*][u] = a[1 : \text{end}]$ 

```

Algorithm 2 InferOpinionParams(Y_u, λ, g_u)

```

1:  $s = \text{numRows}(g_u)$ 
2:  $X = [\vec{1}_s \ g_u]$ 
3:  $Y = Y_u$ 
4:  $L = X'X$ 
5:  $x = (\lambda I + L)^{-1}X'Y$ 
6: return  $x$ 

```

H.2 Estimation of (μ, B)

The first step of the parameter estimation procedure for temporal dynamics also involves the computation of the triggering kernels, which we do in same way as for the opinion dynamics in Algorithm H.1. In order to estimate the parameters, we adopted the spectral projected gradient descent (SPG) method proposed in [4].

Algorithm 3 Spectral projected gradient descent for μ_u and \mathbf{B}_{u^*}

```

1: Input:  $\mathcal{H}_T, G(V, E)$  and  $\mu_u^0, \mathbf{B}_{u^*}^0$ , step-length bound  $\alpha_{max} > 0$  and initial-step length  $\alpha_{bb} \in (0, \alpha_{max}]$ , error bound  $\epsilon$  and
   size of memory  $h$ 
2: Output:  $\hat{\mu}_u, \hat{\mathbf{B}}_{u^*}$ 
3: Notation:
    $x = [\mu_u, \mathbf{B}_{u^*}]$ 
    $f(x) = \sum_{e_i \in \mathcal{H}(T)} \log \lambda_{u_i}^*(t_i) - \sum_{u \in \mathcal{V}} \int_0^T \lambda_u^*(\tau) d\tau$ 
4: Initialize:
    $k = 0$ 
    $x_0 = [\mu_u^0, \mathbf{B}_{u^*}^0]$ 
5: while  $\|d_k\| < \epsilon$  do
6:    $\bar{\alpha}_k \leftarrow \min\{\alpha_{max}, \alpha_{bb}\}$ 
7:    $d_k \leftarrow P_c[x_k - \bar{\alpha}_k \nabla f^u(x_k)] - x_k$ 
8:    $f_b^u \leftarrow \max\{f^u(x_k), f^u(x_{k-1}), \dots, f^u(x_{k-h})\}$ 
9:    $\alpha \leftarrow 1$ 
10:  while  $q_k(x_k + \alpha d_k) > f_b^u + \nu \alpha \nabla f^u(x_k)^T d_k$  do
11:    Select  $\alpha$  by backtracking line-search;
12:     $x_{k+1} \leftarrow x_k + \alpha d_k$ 
13:     $s_k \leftarrow \alpha d_k$ 
14:     $y_k \leftarrow \alpha B_k d_k$ 
15:     $\alpha_{bb} \leftarrow y_k^T y_k / s_k^T y_k$ 
16:     $k = k + 1$ 

```

I Model simulation algorithm

We leverage the simulation algorithm in [10] to design an efficient algorithm for simulating opinion-dynamics. The two basic premises of the simulation algorithm are the sparsity of the network and the Markovian property of the model. Due to sparsity, any sampled event would effect only a few number of intensity functions, only those of the local neighbors of the node. Therefore, to generate the new sample and to identify the intensity functions that require changes, we need $O(\log |\mathcal{V}|)$ operations to maintain the heap for the priority queue. On the other hand, the Markovian property allows us to update the rate and opinion in $O(1)$ operations. The worst-case time-complexity of this algorithm can be found to be $O(d_{\max} |\mathcal{H}(T)| |\mathcal{V}|)$, where d_{\max} is the maximum degree.

Algorithm 4 OpinionModelSimulation(T, μ, B, α, A)

```

1: Initialize the priority queue  $Q$ 
2:  $\text{LastOpinionUpdateTime}[1 : |V|] = \text{LastIntensityUpdateTime}[1 : |V|] = \vec{0}$ 
3:  $\text{LastOpinionUpdateValue}[1 : |V|] = \alpha$ 
4:  $\text{LastIntensityUpdateValue}[1 : |V|] = \mu$ 
5:  $\mathcal{H}(0) \leftarrow \emptyset$ 
6: for  $u \in \mathcal{V}$  do
7:    $t = \text{SampleEvent}(\mu[u], 0, u)$ 
8:    $Q.\text{insert}([t, u])$ 
9: while  $t < T$  do
10:   %% Approve the minimum time of all posts
11:    $[t', u] = Q.\text{ExtractMin}()$ 
12:    $t_u = \text{LastOpinionUpdateTime}[u]$ 
13:    $x_{t_u} = \text{LastOpinionUpdateValue}[u]$ 
14:    $x_{t'_u} \leftarrow \alpha[u] + (x_{t_u} - \alpha[u])e^{-\omega(t - t_u)}$ 
15:    $\text{LastOpinionUpdateTime}[u] = t'$ 
16:    $\text{LastOpinionUpdateValue}[u] = x_{t'_u}$ 
17:    $m_u \sim p(m|x_u(t))$ 
18:    $\mathcal{H}(t') \leftarrow \mathcal{H}(t) \cup (t, m_u, u)$ 
19:   %% Update neighbors'
20:   for  $\forall v$  such that  $u \sim v$  do
21:      $t_v = \text{LastIntensityUpdateTime}[v]$ 
22:      $\lambda_{t_v} = \text{LastIntensityUpdateValue}[v]$ 
23:      $\lambda_{t'_v} \leftarrow \mu[v] + (\lambda_{t_v} - \mu[v])e^{-\omega(t - t_v)} + B_{uv}$ 
24:      $\text{LastIntensityUpdateTime}[v] = t'$ 
25:      $\text{LastIntensityUpdateValue}[v] = \lambda_{t'_v}$ 
26:      $t_v = \text{LastOpinionUpdateTime}[v]$ 
27:      $x_{t_v} = \text{LastOpinionUpdateValue}[v]$ 
28:      $x_{t'_v} \leftarrow \alpha[v] + (x_{t_v} - \alpha[v])e^{-\omega(t - t_v)} + A_{uv}m_u$ 
29:      $\text{LastOpinionUpdateTime}[v] = t'$ 
30:      $\text{LastOpinionUpdateValue}[v] = x_{t'_v}$ 
31:     %% Sample by only effected nodes
32:      $t_+ = \text{SampleEvent}(\lambda_v(t'), t', v)$ 
33:      $Q.\text{UpdateKey}(v, t_+)$ 
34:    $t \leftarrow t'$ 
35: return  $\mathcal{H}(T)$ 

```

Algorithm 5 SampleEvent(λ, t, v)

```

1:  $\bar{\lambda} = \lambda, s \leftarrow t$ 
2: while  $s < T$  do
3:    $\mathcal{U} \sim \text{Uniform}[0, 1]$ 
4:    $s = s - \frac{\ln \mathcal{U}}{\bar{\lambda}}$ 
5:    $\lambda(s) \leftarrow \mu[v] + (\lambda_v(t) - \mu[v])e^{-\omega(s - t)}$ 
6:    $d \sim \text{Uniform}[0, 1]$ 
7:   if  $d\bar{\lambda} < \lambda$  then
8:     break
9:   else
10:     $\bar{\lambda} = \lambda(s)$ 
11: return  $s$ 

```

J Experiments on Synthetic Data

Parameter estimation accuracy. We evaluate the accuracy of our model estimation procedure on three types of Kronecker networks [19]: i) assortative networks (parameter matrix $[0.96, 0.3; 0.3, 0.96]$), in which nodes tend to link to nodes with similar degree; ii) disassortative networks ($[0.3, 0.96; 0.96, 0.3]$), in which nodes tend to link to nodes with different degree; and iii) core-periphery networks ($[0.9, 0.5; 0.5, 0.3]$). For each network, the message intensities are multivariate Hawkes, μ and B are drawn from a uniform distribution $U(0, 1)$, and α and A are drawn from a Gaussian distribution $\mathcal{N}(\mu = 0, \sigma = 1)$. We use exponential kernels with parameters $\omega = 100$ and $\nu = 1$, respectively, for opinions $\mathbf{x}_u^*(t)$ and intensities $\lambda^*(t)$. We evaluate the accuracy of our estimation procedure in terms of mean squared error (MSE), between the estimated and true parameters, *i.e.*, $\mathbb{E}[(x - \hat{x})^2]$. Figure 5 shows the MSE of the parameters (α, A) , which control the Hawkes intensities, and the parameters (μ, B) , which control the opinion updates. As we feed more messages into the estimation procedure, it becomes more accurate.

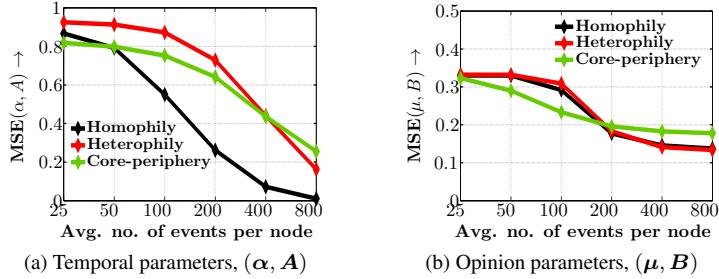


Figure 5: Performance of model estimation for several 512-node kronecker networks in terms of mean squared error between estimated and true parameters. As we feed more messages into the estimation procedure, the estimation becomes more accurate.

K Twitter dataset description

We used the Twitter search API⁸ to collect all the tweets (corresponding to a 2-3 weeks period around the event date) that contain hashtags related to the following events/topics:

- **Politics:** Delhi Assembly election, from 9th to 15th of December 2013.
- **Movie:** Release of “*Avengers: Age of Ultron*” movie, from April 28th to May 5th, 2015.
- **Fight:** Professional boxing match between the eight-division world champion Manny Pacquiao and the undefeated five-division world champion Floyd Mayweather Jr., on May 2, 2015.
- **Bollywood:** Verdict that declared guilty to Salman Khan (a popular Bollywood movie star) for causing death of a person by rash and negligible driving, from May 5th to 16th, 2015.
- **US:** Presidential election in the United-States, from April 7th to 13th, 2016.

We then built the follower-followee network for the users that posted the collected tweets using the Twitter rest API⁹. Finally, we filtered out users that posted less than 200 tweets during the account lifetime, follow less than 100 users, or have less than 50 followers.

Dataset	$ \mathcal{V} $	$ \mathcal{E} $	$ \mathcal{H}(T) $	$\mathbb{E}[\mathbf{m}]$	$\text{std}[\mathbf{m}]$
Tw: Politics	548	5271	20026	0.0169	0.1780
Tw: Movie	567	4886	14016	0.5969	0.1358
Tw: Fight	848	10118	21526	-0.0123	0.2577
Tw: Bollywood	1031	34952	46845	0.5101	0.2310
Tw: US	533	20067	18704	-0.0186	0.7135

Table 1: Real datasets statistics

⁸<https://dev.twitter.com/rest/public/search>

⁹<https://dev.twitter.com/rest/public>